

CLAIMS

What is claimed is:

1. A system that facilitates query optimization in a data repository, comprising:
 - a query component that receives a query to be processed against data of the data repository, which query includes an original predicate; and
 - a predicate component that processes the original predicate into one or more new predicates that include an implied predicate, the implied predicate processed against the data to return a best solution such that a total evaluation cost is significantly reduced.
2. The system of claim 1, the predicate component processes the original predicate to obtain an equivalent predicate.
3. The system of claim 1, the predicate component processes the original predicate to obtain a residual predicate.
4. The system of claim 1, the query is a previously processed query that is unmatched.
5. The system of claim 1, the one or more new predicates can be used for standard indices and multi-valued indices.
6. The system of claim 1, the one or more new predicates are considered cost-based alternatives that are utilized only if the evaluation cost is reduced, otherwise, they are discarded.
7. The system of claim 1, the one or more new predicates are associated with index utilization.

8. The system of claim 1, the query is processed against a data type that is a non-indexable type.

9. A data repository optimizer according to the system of claim 1.

10. A system that facilitates query optimization in a data repository, comprising:

a query component that receives a query to be processed against data of the data repository, which query includes an original predicate; and

a predicate component that processes the original predicate into one or more new predicates that include at least one of an implied predicate, an equivalent predicate, and a residual predicate, either of the implied predicate or the equivalent predicate is processed against the data to return a best solution such that a total evaluation cost is significantly reduced.

11. The system of claim 10, the implied predicate and the equivalent predicate are considered cost-based alternatives that are discarded if the evaluation cost is not reduced.

12. A computer readable medium having stored thereon computer executable instructions for carrying out the system of claim 10.

13. A server that employs the system of claim 10.

14. The system of claim 10, the predicate component generates an expression using the one or more new predicates, which expression is used to obtain an implication rule that is associated with at least one of a given column and a function on a given column.

15. The system of claim 14, the predicate component requests the implication rule in response to providing the expression to which at least one of the given column and the function on a given column is to be compared, and a list of standard indices that can be exploited.

16. The system of claim 10, the one or more new predicates is analyzed during index selection.

17. The system of claim 10, further comprising a classifier that facilitates automating one or more feature thereof by makes an inference based on one or more parameters related to at least one of a cost-basis evaluation, cardinality estimation, and complexity of the query.

18. The system of claim 10, the total evaluation cost includes employing the one or more new predicates only if optimization is increased.

19. The system of claim 10, the query is one for which there is no exact match between search predicates and index keys.

20. The system of claim 10, the implied predicate exactly matches an index key.

21. A computer-readable medium having computer-executable instructions for performing a method for optimizing a search query, the method comprising:

receiving a query for whose original predicate there is no exact match to an index key;

reducing the original predicate into at least one of an implied predicate and an equivalent predicate;

processing at least one of the implied predicate and the equivalent predicate against data of a data repository to obtain search results; and

analyzing the search results for a best solution.

22. The method of claim 21, further comprising requesting an implication rule for a column or function on a column.
23. The method of claim 21, further comprising transmitting at least one of an expression to which a column or function on a column is to be compared and a list of standard indices or multi-valued indices that could be exploited.
24. The method of claim 21, further comprising removing the implied predicate if the implied predicate does not exploit one of the standard or multi-value indices.
25. The method of claim 21, further comprising employing a new ad-hoc exploratory rule for an individual operator of the original predicate.
26. The method of claim 25, further comprising analyzing at least one of the implied predicate and the equivalent predicate with the exploratory rule.
27. The method of claim 21, further comprising the acts of:
requesting a list of tentative substitutes with predicates that involve nested queries;
including the substitutes into the query to form a new expression;
simplifying the new expression;
removing nested queries; and
generating new alternatives for the search.
28. The method of claim 21, further comprising processing the original predicate to generate a residual predicate the refines the search.

29. The method of claim 28, further comprising processing the residual predicate after the act of processing at least one of the implied predicate and the equivalent predicate.

30. The method of claim 21, further comprising adding the implied predicate to the original predicate and searching for the best solution.

31. The method of claim 21, further comprising replacing the original predicate with the equivalent predicate and searching for the best solution.

32. The method of claim 21, further comprising performing a cardinality estimation using at least one of the implied predicate and the equivalent predicate.

33. The method of claim 21, further comprising creating an index based on an index extension scheme, which scheme includes at least one of,

providing a key column of a data type that corresponds to a data type being indexed;

providing a set of parameters; and

providing a table-valued function name that is used to generate an index entry for a value of the column being indexed.

34. A system that facilitates query optimization in a data repository; comprising:

- means for receiving a query for whose original predicate there is no exact match to an index key;
- means for reducing the original predicate into at least one of an implied predicate and an equivalent predicate;
- means for requesting an implication rule for a column or function on a column;
- means for transmitting at least one of an expression to which a column or function on a column is to be compared and a list of standard indices or multi-valued indices that could be exploited.
- means for processing at least one of the implied predicate and the equivalent predicate against data of a data repository to obtain search results;
- means for analyzing the search results for a best solution; and
- means for removing the implied predicate if the implied predicate does not exploit one of the standard or multi-value indices.

35. The system of claim 34, further comprising means for employing a new ad-hoc exploratory rule for an individual operator of the original predicate.

36. The system of claim 34, further comprising means for analyzing at least one of the implied predicate and the equivalent predicate with the exploratory rule.

37. The system of claim 34, further comprising at least one of:

- means for requesting a list of tentative substitutes with predicates that involve nested queries;
- means for including the substitutes into the query to form a new expression;
- means for simplifying the new expression;
- means for removing nested queries; and
- means for generating new alternatives for the search.

38. The system of claim 34, further comprising means for processing the original predicate to generate a residual predicate the refines the search.

39. The system of claim 38, further comprising means for processing the residual predicate after processing at least one of the implied predicate and the equivalent predicate.

40. The system of claim 34, further comprising means for estimating cardinality using at least one of the implied predicate and the equivalent predicate, where use of the implied predicate facilitates defining an upper bound on a number of rows that can satisfy the original predicate.